

SRI International



SHAPE FROM SHADING: AN ASSESSMENT

Technical Note 287

May 1983

By: Grahame B. Smith

Artificial Intelligence Center
Computer Science and Technology Division

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE MAY 1983		2. REPORT TYPE		3. DATES COVERED 00-05-1983 to 00-05-1983	
4. TITLE AND SUBTITLE Shape from Shading: An Assessment				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) SRI International, 333 Ravenswood Avenue, Menlo Park, CA, 94025				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 20	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

Abstract

We review previous efforts to recover surface shape from image irradiance in order to assess what can and cannot be accomplished. We consider the informational requirements and restrictions of these approaches. In dealing with the question of what surface parameters can be recovered locally from image shading, we show that, at most, shading determines relative surface curvature, i.e., the ratio of surface curvature measured in orthogonal image directions. The relationship between relative surface curvature and the second derivatives of image irradiance is independent of other scene parameters, but insufficient to determine surface shape. This result places in perspective the difficulty encountered in previous attempts to recover surface orientation from image shading.

1. Introduction

The determination of land cover from aerial imagery is a task that photo interpreters accomplish by using both the image data and their knowledge of the structure of the world. The image data encodes the complex process whereby light is reflected from a surface. The surface shape, the surface albedo, the position of the lighting sources, and the functional form of the reflectance properties of the material are elements of this encoding. The human visual system interprets image data as a 3-D model of the scene, distinguishes among different surface materials, and ascertains the position of the lighting sources. It is difficult to believe that a machine vision system can achieve, say, surface material differentiation without simultaneously being able to recover the surface shape and the other parameters that are needed to explain the detected image intensity. Of course, it may be possible to use special sensors and multiple information sources to make it unnecessary to reconstruct a complete 3-D model of the scene, but it would be surprising if such specialization could retain sufficient generality to be useful over a range of remote sensing tasks, e.g., in both renewable and nonrenewable resources.

The machine vision approach of simultaneously recovering all the parameters necessary to account for image intensity is expressed in the notion of intrinsic images [1] (or the $2\frac{1}{2}$ -D sketch [2]). These intrinsic images can be thought of as overlays, each specifying the value of one parameter that goes into the formula for calculating the image intensity. The

The research reported herein was supported by the Defense Advanced Research Projects Agency under Contract MDA903-83-C-0027 and by the National Aeronautics and Space Administration under Contract NASA 9-16864. These contracts are monitored by the U.S. Army Engineer Topographic Laboratory and by the Texas A&M Research Foundation for the Lyndon B. Johnson Space Center.

images are not independent; if one is to be varied, the others must be also — so that the predicted image intensity remains invariant (and equal to the observed value). The notional division of an image into particular intrinsic images would be of little merit unless one believed that estimates of each intrinsic image could be obtained by models that were largely independent of the other intrinsic images. While models have been proposed to recover various intrinsic images, there have been considerable efforts made to recover the scene's 3-D shape,¹ in particular the surface orientation at each image point. These 'shape-from-...' models embody a structure that would allow shape to be recovered principally from a single measure, e.g., texture, contour, or shading. While 'shape-from-...' models are not seen as complete solutions to shape reconstruction, there is an implicit expectation in their title that shape estimates can be calculated from their respective measures. Here we review the work we and others have done towards the goal of recovering surface shape from image shading. Is it attainable — or is it myth?

The importance of shape recovery is clear; if the shape is known, surface albedo, and the other parameters that determine image intensity are obtainable. Land cover differentiation is dependent on knowing the [relative] surface albedo, rather than image parameters, such as intensity. If we cannot recover shape, the intrinsic image approach offers little as a model for perception. Shading is only one source of shape information. Edge information is of great importance, but there is little occlusion in aerial images. The ability to recover shape from shading seems more critical in the case of aerial imagery than for most other types of imagery.

We first review three research efforts: those of Horn and his colleagues [3-8], Pentland [9], and our own [10], — to determine what can and cannot be accomplished, and to consider the informational requirements and limitations of these approaches. We discuss the dilemma of local computation versus global constraint propagation and seek to ascertain what can be computed locally, and how information can be propagated across an image. Finally, we seem to be left with the conclusion that shading, when viewed as a single source of shape information, is an insufficient source for the recovery of surface shape. Shape cannot be obtained from shading alone. However, we are able to characterize the scene information that shading provides.

An alternative approach to recovering shape from shading is model based. Can we determine which model, from a set of models, best describes the image data? This approach is dependent on discovering a small set of easily distinguishable models that adequately describe the surfaces encountered. Industrial inspection, rather than remote sensing of the environment, appears better suited to a model based procedure. In this assessment we do not consider this related, but essentially different approach.

¹We use the expression surface shape to denote both the intrinsic properties of the surface, e.g., cylindrical, and the orientation of the surface in space. Elsewhere, shape is sometimes used to denote only the intrinsic properties of the surface, not its orientation in space.

2. Approaches to Shape from Shading

2.1 Horn and Colleagues

A study by Horn [3,4,11] of the relationship among image irradiance,² surface shape, surface albedo, and illumination conditions led to formulation of the image irradiance equation, which states that image irradiance is proportional to scene radiance.³ This is expressed by the equation

$$I = R \quad ,$$

where I is the image irradiance as a function of the image coordinates, and R is the scene radiance as a function of the scene parameters. Of course, this equation relates the image irradiance at an position in the image to the scene radiance at its corresponding scene position. Implicit in this equation is an assumption of orthographic projection. However, such an assumption, to avoid complexity in the mathematical formulation, is a minor restriction and does not detract from the generality of the model.

Image irradiance is a function of the image coordinates x and y , but scene radiance is a function of the illumination strength, its position, the surface albedo, and the surface orientation. For the formulations reviewed here, we find that a number of assumptions are made so that scene radiance can be considered a function of the surface orientation variables only; constant values are used for the illumination strength, its position, and for the surface albedo. That is, shape-from-shading is analyzed for the simplified case of a constant light source and constant surface albedo. The restriction to a constant light source is not only a good approximation of the situation we experience daily (and an excellent approximation for a photograph), but also corresponds to the difficulty confronting the human visual system when this constancy is not met, e.g., under strobe lighting. The assumption of constant albedo is harder to justify, since nature obviously exhibits variable albedo. Still, when we consider the manner in which facial make-up is used to alter the perceived shape of the face, it may well be that continuous changes in albedo are processed by the human visual system as if they were constant. Notwithstanding the justification for constant albedo, it is unlikely that shape-from-shading can be solved for the case of variable albedo if it cannot be solved for constant albedo. Such a restriction is in effect a case analysis to determine if shading provides sufficient shape information in a less-than-general model.

In the formulations under review, various parameterizations of surface orientation have been used. The two we specify are (i) surface gradients, i.e., the partial derivatives of depth, z , with respect to the scene (and image) coordinates x and y , and (ii) components of the surface normal, i.e., l and m , the x and y components of the surface normal. Using the notation, $p = \frac{\partial z}{\partial x}$, and $q = \frac{\partial z}{\partial y}$, we note the equivalence of the parameterizations

$$p = \frac{-l}{\sqrt{1-l^2-m^2}} \quad , \text{ and } \quad q = \frac{-m}{\sqrt{1-l^2-m^2}} \quad .$$

²Image irradiance is the light flux per unit area falling on the image, i.e., incident flux density.

³Scene radiance is the light flux per unit projected area per unit solid angle emitted from the scene, i.e., emitted flux density per unit solid angle.

The image irradiance equation is usually expressed as

$$I(x, y) = R(p, q) \quad , \text{ or } \quad I(x, y) = R(l, m) \quad ,$$

and we shall use both forms to express the relationship between image irradiance and scene radiance for the case of constant illumination and constant albedo. As $p = \frac{\partial z}{\partial x}$, and $q = \frac{\partial z}{\partial y}$, we see that the image irradiance equation is a first-order partial differential equation and, if I and R are known, we could (at least in principle) solve the differential equation and recover the depth, z .

To have an explicit form for R , we must have a model for the type of reflection occurring at the scene surfaces. In the work reviewed here the surface is assumed to be a perfectly uniform diffuse reflector, i.e., the scene radiance is isotropic.⁴ While this model is invalid as a description of specular reflection, scene radiance in the natural world, (except for specific situations, such as water surfaces), may be approximated by such a description. The expression for scene radiance in this case is [10]

$$R(l, m) = al + bm + c\sqrt{1 - l^2 - m^2}$$

or, equivalently,

$$R(p, q) = \frac{(-ap - bq + c)}{\sqrt{1 + p^2 + q^2}} \quad ,$$

where a, b , and c are constants expressing illumination strength, its position, and the surface albedo.

The approach taken by Horn and his colleagues [3-8] is to solve the first-order partial differential equation,

$$I(x, y) = \frac{(-ap - bq + c)}{\sqrt{1 + p^2 + q^2}} \quad ,$$

assuming that a, b , and c are known — i.e., the surface albedo, and the illumination strength, and its position. While this need to know scene parameters may seem over-restrictive, such information may come from other components of a vision system. The need to know the illumination position does not seem to be a major drawback of this approach, but the requirement that the scene albedo be known is troublesome. If the conceptual model of intrinsic images is to be followed, the inability to decouple surface orientation from surface albedo would seem fundamental. Regardless of this difficulty, the question of whether shape can be recovered in a limited domain is basic to the investigation of vision.

Two approaches to solving the image irradiance equation are direct integration [3,4], and iterative/relaxation techniques [5-8]. The direct integration approach has been carried out generally in those circumstances in which $I(x, y)$ and its derivatives can be determined for all x and y , i.e., for a spatially unquantized, continuous-tone image. The method used is the standard technique of characteristic strips for solving a first-order hyperbolic partial

⁴This situation is also called Lambertian reflectance, after Lambert, who proposed a point reflection model (in which the reflected flux per unit surface area per unit solid angle varied as the cosine of the angle between the surface normal and the viewing direction) to account for the observation that matt surfaces looked equally bright from any viewing position.

differential equation [3,4]. Starting with a point at which the surface orientation is known, integration moves along a curve in the image. This curve is dictated by the image. Adjacent curves generally are not 'parallel', which makes it difficult to get complete coverage of the image. Interpolation between these curves — or strips, as they are usually called — to find initial values to commence an intervening strip integration, involves complex procedures. As far as digital images are concerned, direct integration would be hard to organize, even if we were first to model the intensities to obtain a continuous form for $I(x, y)$.

As is the case with most partial differential equations, it should be noted that the image irradiance equation has many solutions [12]. The boundary conditions (in the above method the initial values for a strip) are vital in selecting the solution that describes the surface in the image. Should the image irradiance equation be 'underconstrained' in the sense that, for a given $I(x, y)$, it admits solutions that encompass a wide range of surface types with similar boundary values, we might then expect numerical error to defeat attempts at numerical integration. In such cases, errors 'mix in' other solutions that can eventually dominate the recovered solution, even though they may be excluded by the boundary conditions. The method of direct integration has been demonstrated on simple images [3]. These examples required only a small number of integration steps. Numerical instability has also been reported [12].

The other approach used to solve the image irradiance equation is relaxation. Relaxation procedures avoid numerical instability, but face the problem of convergence. However, they do have the advantage of being directly applicable to digital images, i.e., spatially quantized, discrete-tone images. The relaxation (or iterative) approach views the image irradiance equation not as a differential equation, but as an algebraic constraint. For pixel (i, j) ,

$$I_{i,j} = R(p_{i,j}, q_{i,j}) \quad ,$$

where $I_{i,j}$ is the image irradiance for the (i, j) th pixel, and $p_{i,j}$ and $q_{i,j}$ specify the surface orientation of the surface patch that is imaged at pixel (i, j) . As an algebraic constraint, the image irradiance equation relates image irradiance to the two surface orientation variables, $p_{i,j}$ and $q_{i,j}$. In viewing the image irradiance equation as an algebraic constraint, we lose the interrelationship of $p_{i,j}$, $q_{i,j}$, and their neighboring values, a relationship inherent in the differential equation. To compensate for this loss, an additional constraint must be introduced that relates $p_{i,j}$ and $q_{i,j}$ to their neighboring values. Such a relationship is essential for a relaxation procedure. The relationship usually introduced attempts to capture the notion of surface smoothness [7,8, 10,13]. The particular form of the smoothness constraint may, for example, require that $p_{i,j}$ and $q_{i,j}$ be equal to the mean values of neighboring p 's and q 's. For any trial values for $p_{i,j}$ and $q_{i,j}$, the constraint imposed by the image irradiance equation and the constraint resulting from surface smoothness will not be completely satisfied. The residual equation formed from each constraint specifies how well that constraint is satisfied. If $\xi_{i,j}$ is the sum of the [absolute values of the] residuals from both the image irradiance constraint and the surface smoothness constraint for the (i, j) th pixel, then, for trial values of p and q for every image pixel, the total residual error is

$$\xi = \sum_{i,j \in \text{image}} \xi_{i,j} \quad .$$

The allocation of surface orientations to all pixels should minimize this total error — that is,

$$\begin{aligned}\frac{\partial \xi}{\partial p_{i,j}} &= 0 & \forall i,j \in \text{image} \quad , \\ \frac{\partial \xi}{\partial q_{i,j}} &= 0 & \forall i,j \in \text{image} \quad .\end{aligned}$$

From these equations we obtain an iterative scheme for updating the values of p and q so that they are compatible with their neighboring values, as well as with the image irradiance equation [7,8,10]. If such a scheme is convergent, we have a procedure for obtaining shape from shading.

It should be noted that the relaxation schemes, that use the foregoing approach are possible only because the smoothness constraint relates the values at one pixel to those of its neighboring pixels. The boundary conditions needed for selecting a particular solution from the solution set of the iterative scheme are propagated by the smoothness constraint, not the image irradiance equation. Compared with the direct-integration approach, information propagation in the relaxation scheme uses a different mechanism. We must remember this when we assess results.

Success with these methods has generally been limited to small images, (usually fewer than 30 x 30 pixels), of nearly spherical or saddle surfaces [7,8,10,14]. For an effective relaxation scheme, the initial solution should have no effect on the surface recovered. This unfortunately is not the case [10]. Boundary conditions are not propagated more than a few pixels by the smoothness constraints [7,10]. Surface recovery from large images, (bigger than 30 x 30 pixels), is ineffectual for this reason. As a consequence of the fact that smoothness is used as information propagator, assumptions (albeit weak ones) have been made about surface shape. Shading as a constraint, and smoothness as a surface type, appears insufficient to provide a basis for an effective shape-from-shading algorithm.

2.2 Pentland

The approaches to solving the shape-from-shading task that we have discussed so far have all been based on constraint propagation. Direct integration is a spatially serial solution to the propagation problem, while relaxation attempts to achieve this propagation with a temporally serial solution; in other words, relaxation employs local processing, but it must iterate until enough cycles have passed to allow information to propagate spatially. Purely local computation of scene parameters, on the other hand, is not a propagation method. While this kind of computation can use neighboring data — and not just of the nearest neighbors — it must provide an instant solution. It cannot iterate and therefore it does not provide a temporally serial solution. Such an approach to scene parameter computation avoids the numerical instability of direct integration methods, as well as the convergence and propagation problems of relaxation, but it cannot use spatially distant scene information. A local computation can use global information, such as the position of the light source, but it cannot use scene details, such as the position of a distant edge. Of course, the reason for carrying out purely local computation stems from the hypothesis that such scene detail is not involved in the computation at this level in the visual system. Can shading provide

sufficient local information to allow recovery of surface shape by purely local computation? This is the question addressed by Pentland [9].

The inadequacy of local image measurements for specifying surface orientation can be understood by counting the variables needed to specify various image measurements. Let us consider the case of a uniformly diffuse reflecting surface. Image irradiance (1 measurement) is a function of surface orientation (2 parameters), the product of surface albedo and illumination strength (1 parameter), and the position of the light source (2 parameters). The gradients of image irradiance (2 measurements) are functions of the same variables as image irradiance and, additionally, are functions of surface curvature (3 parameters). The second derivatives of image irradiance (3 measurements) are functions of all the variables mentioned above, plus the rates of change of curvature (4 parameters). Because higher image-irradiance derivatives introduce more surface shape derivatives, we have more parameters than measurements. It should be noted that a knowledge of global quantities, such as the illumination position and the product of surface albedo and illumination strength, is not sufficient to allow the surface orientation to be computed locally. If we make assumptions about the relationship among some of the above parameters, we can produce a system of equations from which surface orientation can be calculated.

Pentland investigates the case in which an image patch of a uniformly diffuse reflecting surface can be considered identical to a point on an illuminated sphere whose reflection is also uniformly diffuse [9]. He calculates the orientation of the surface patch on the sphere that has the same appearance as the surface patch in the image. Not all image patches can be represented by points on an illuminated sphere. Spheres whose reflection is uniformly diffuse have the property

$$\frac{I_{xx}}{I_{yy}} \geq 0 \quad ,$$

where subscripts denote partial differentiation with respect to those subscripts. There are surfaces, e.g., a sinusoidal surface, for which $\frac{I_{xx}}{I_{yy}}$ can be negative. The procedure for estimating surface orientation that is based on the assumption that surfaces can be approximated by locally spherical patches is applicable only to parts of an image. Notwithstanding these restrictions, an important aspect of the assumption of local sphericity is that the surface orientation is calculated by using the second derivatives of image irradiance only, i.e.,

$$\begin{aligned} \frac{1 - m^2}{lm} &= \frac{I_{xx}}{I_{xy}} \quad , \\ \frac{1 - l^2}{lm} &= \frac{I_{yy}}{I_{xy}} \quad . \end{aligned}$$

These equations are derived by differentiating the image irradiance equation and noting that, for a sphere, $l_x = \frac{1}{r}$, $l_y = 0$, $m_x = 0$, and $m_y = \frac{1}{r}$, where r is the sphere's radius.

In this model, surface orientation is directly dependent on neither image irradiance nor on the first derivatives of image irradiance. It may be estimated even in images that exhibit linear changes in irradiance induced by artifacts, and in images that exhibit constant illumination levels induced by atmospheric effects, such as backscatter. More importantly, the formulas are independent of the illumination parameters and the surface albedo. In exchange for acceptance of a restrictive assumption with respect to surface type,

one acquires not only a means of calculating surface orientation, but a procedure that needs no information other than image measurements — a procedure, in effect, that is matched to the notion of intrinsic images.

Even in those areas of an image to which this approximation can be applied, the assumption that a surface can be approximated by a patch with the same curvature in any direction needs experimental verification. The world is obviously not composed of such surfaces, but it is the difference between the estimated and the actual surface orientation that is more important than the error made in approximating the surface by a spherical patch. Application of the above formula yields qualitative agreement between the estimated and actual shape in synthetic images and in natural images of simple objects [9], (for which $\frac{I_{xx}}{I_{yy}}$ is generally positive). Shape estimates from synthetic images of ellipsoidal surfaces are 'flatter' than the actual shapes. It should be noted that shape estimates, which are integrated surface orientations, often appear 'better' than what might be expected on the basis of the surface orientation error. An algorithm based on approximating a surface patch by a spherical one seems better suited for computing the qualitative shape of a surface than the orientation of surface elements. Such an algorithm is applicable only to those image patches that are consistent with the interpretation of such patches as points on a sphere. The conditions necessary for enabling this kind of interpretation have not been fully characterized. Alternative models, that are applicable when an image patch is inconsistent with an interpretation that it is a point on a sphere, are currently unknown.

In principle, because image irradiance is not differentiable at boundaries, we cannot apply the above method there. However, unlike propagation methods require our knowing boundary positions in order to stop computation, the local-computation approach may accomplish this simply by indicating (through its failure at a boundary) where the boundary is.

Pentland's approach hinges on the local-sphericity assumption. In restricted circumstances he is able to estimate surface orientation directly from the second derivatives of the image irradiance. What other, perhaps less specific, assumptions can be made that allow shape to be estimated locally? Before attempting to answer this, we review the shape-from-shading formulation we have previously proposed [10,15], — first, to assess its performance, then to provide the requisite analytical tools for answering questions about local computation.

2.3 Smith

The approach taken by Horn and his colleagues provides a formulation of the shape-from-shading task that requires knowledge of scene parameters, but places no restriction on the surface shape. Calculation of surface orientation is not a local process, and, if surface orientation is to be recovered, knowledge of boundary conditions is necessary. Pentland, on the other hand, restricts the surface shape but requires no scene parameters, no boundary conditions, and derives surface orientation by purely local computation. Is there an intermediate position? Is there a formulation that neither restricts the surface shape nor requires knowledge of scene parameters? Of course, local computation seems desirable — but is it worth the concomitant cost of surface type restriction or the requirement that scene parameters be known a priori? The formulation previously described by us, takes such an

intermediate position.

For a uniformly diffuse reflecting surface, surface orientation is related to image irradiance by the second-order partial differential equations [10]

$$\begin{aligned}\alpha\theta l_{xx} + \beta\theta m_{xx} - \alpha\gamma l_{xy} - \beta\gamma m_{xy} &= \chi\theta I_{xx} - \chi\gamma I_{xy} \quad , \\ \alpha\theta l_{yy} + \beta\theta m_{yy} - \alpha\delta l_{xy} - \beta\delta m_{xy} &= \chi\theta I_{yy} - \chi\delta I_{xy} \quad ,\end{aligned}$$

where

$$\begin{aligned}\alpha &= I_x m_y - I_y m_x \quad , \\ \beta &= I_y l_x - I_x l_y \quad , \\ \gamma &= l_x^2(1 - m^2) + m_x^2(1 - l^2) + 2l_x m_x l m \quad , \\ \delta &= l_y^2(1 - m^2) + m_y^2(1 - l^2) + 2l_y m_y l m \quad , \\ \theta &= l_x l_y(1 - m^2) + m_x m_y(1 - l^2) + (l_x m_y + l_y m_x) l m \quad , \\ \chi &= l_x m_y - l_y m_x \quad .\end{aligned}$$

These equations are derived from the image irradiance equation. The assumption of uniformly diffuse reflection relates some of the scene parameters, thereby allowing elimination of parameters that specify surface albedo and illumination conditions.

The assumption that surface reflection is uniformly diffuse is an assumption about the physics of image formation. While it does not describe the reflectance properties of all surface, it is a reasonable approximation to most surfaces that are encountered in the natural world. For any formulation of the relationship between shading and shape, some assumptions are necessary. Those describing properties found in nature are more palatable than restrictions for which little a priori evidence is available.

A desirable aspect of this formulation is that surface orientation is not related to image irradiance, but only to its derivatives. The existence of constant illumination levels, from atmospheric scattering or fogging of photographic images, does not impede the potential for shape recovery. Linear changes in intensity, however, must affect the shape of any recovered surface. A more important aspect of this formulation is its independence of surface albedo. Again we reiterate that, if the notion of intrinsic images is to be useful we must find models that decouple surface shape from surface reflectance. The fact that knowledge of the illumination conditions is not required, is certainly an important aspect, but less so than the formulation's independence of surface albedo.

The penalty for not making assumptions about surface type and for not presupposing any knowledge of scene parameters, such as illumination conditions and surface albedo, is the introduction of higher-order derivatives of surface orientation in the formulation, as well as the inability to calculate surface orientation by purely local computation. Boundary conditions are necessary. To formulate a model that relates surface orientation to image irradiance is one thing; to solve it for that orientation is another.

The second-order partial differential equations (given above) relating surface orientation and image irradiance are satisfied by solutions to the first-order partial differential equation $\chi = 0$. This is undesirable, as solutions of $\chi = 0$ satisfy the surface-orientation-image-irradiance equations independently of the image measurements, I_x, I_y, I_{xx}, I_{yy} , and I_{xy} . The equation $\chi = 0$ characterizes the developable surfaces, e.g., a cylinder or cone (see Appendix B); derivation of the above surface-orientation-image-irradiance equations is impossible when

the surface is developable, i.e., singularly curved. The surface-orientation-image-irradiance equations are appropriate only when the surface is doubly curved. For singularly curved surfaces, the appropriate equations relating surface orientation and image irradiance are

$$I_x^2(m_y(1 - l^2) + l_y lm) = I_y^2(l_x(1 - m^2) + m_x lm) \quad ,$$

$$I_x m_y - I_y m_x = 0 \quad ,$$

$$I_x l_y - I_y l_x = 0 \quad .$$

(These equations are derived independently of any reflection function, i.e., they apply to all types of reflection, not just uniformly diffuse reflection. See Appendix C.)

If the surface-orientation-image-irradiance equations were solved by analytic procedures, the problems posed by the $\chi = 0$ solutions would vanish, as such solutions would be ruled out by boundary conditions. However, the presence of such solutions heralds difficulties for numerical methods, as the inevitable numerical errors will mix these solutions into the recovered surface orientations. Two approaches to solving the surface-orientation-image-irradiance equations have been reported [15]. These approaches are direct integration, which is implemented by finite-difference formulas, and relaxation. Both require additional information in the form of boundary conditions. Both fail to recover surface orientation. Direct integration correctly recovers the surface orientation in the vicinity of the boundary conditions, but is ineffective elsewhere. The reasons for failure of each method are of interest; direct integration fails because numerical instability makes the spatially serial method of solution impractical; relaxation fails because nonconvergence makes the temporally serial method of solution infeasible. These direct reasons for failure mask a deeper problem. The model is 'underconstrained' from the standpoint that the equations are insensitive to surface orientation. They are more sensitive to other surface parameters, such as surface curvature [15]. Underconstraint of the model can account for lack of convergence of relaxation methods, but the numerical problems in direct integration highlight the difficulty of spatial information propagation by a mechanism that is under the control of higher-order derivatives.

The surface-orientation-image-irradiance equations alone do not form the basis for an algorithm to recover surface orientation; they do provide a tool, however, for examining the constraint shading imposes on shape. We shall subsequently use them for that purpose.

3. Local Computation Versus Global-Constraint Propagation

What can we learn from these various approaches to shape-from-shading? Direct integration of a differential model is an inadequate computational tool. Horn and his colleagues, using a low-order partial differential equation, show that some propagation of information is possible — but numerical instability poses a difficulty even for a first-order equation. This limited success with direct integration is unlikely to be upgradable to a solution procedure for natural scenes. Since higher-order formulations are plagued with numerical instability they do not offer any prospect of success.

A restricting factor in a differential model is the need for knowledge of boundary conditions. This seems to be a major limitation of such methods. These methods apply

to continuous surface patches only and require a priori knowledge of solution values at some points within every region. This means that we must find regional boundaries — perhaps ascertain their type and estimate values of surface orientation at some points within each region before we can attempt to recover shape. Is this, in effect, placing the cart before the horse?

Models of the relationship between image measurements and scene variables that are formulated as low-order differential equations offer no relief from the necessity of knowing scene parameters. While information about illumination conditions may be obtainable from other sources within the image, or maybe calculated in parallel with shape, it is difficult to envisage a situation in which the surface albedo could be calculated before the surface shape. Albedo would seem less constrained than shape. The author's higher-order differential equations show that derivatives of image irradiance can be used to remove these parameters.

While the relaxation schemes used to solve the image irradiance equation are not quite viable, their drawbacks may be attributed to the weakness of the surface shape constraint, namely smoothness, rather than an inherent deficiency of relaxation as a technique. For the higher-order surface-orientation-image-irradiance equations, insensitiveness to surface orientation does not allow assessment of the strength of surface continuity (the constraint used in the attempts to solve these equations by relaxation). The results reported from these relaxation procedures can be attributed to other aspects of the models they embody, rather than to any deficiency of the relaxation technique itself. Relaxation seems viable as a method that can satisfy global constraints without being dominated by numerical error. However, surface shape assumptions, that are more restrictive than those used in the work reviewed, appear necessary if information is to be propagated effectively over reasonable image distances. Relaxation schemes that implement low-order differential models seem practicable; schemes implementing higher-order differential models are too sensitive to noise.

In comparison with information propagation techniques, local computation of surface orientation, as reported by Pentland, requires strong restrictions on surface shape — and even these are not adequate to characterize all cases. However, local computation, particularly when it is based on a model involving derivatives of image irradiance only, does provide a means for recovering surface orientation without any knowledge of boundary conditions, without a priori regional segmentation (it may even help in this endeavor), and without knowing the scene parameters, especially albedo. Unfortunately, we shall not get a solution to surface orientation that is quantitatively correct because the surface restriction is too great. Local computation offers the computational features we want, but the penalty to be paid — severe surface shape restriction — is far too great.

What, then, seems practical? A relaxation scheme that is more constrained by surface type than those that have been examined? A scheme that implements a low-order model of information propagation? A scheme that does a lot of purely local computation? A scheme that can use boundary conditions wherever they are, but without being overly dependent on them? Of course, all this is one conjecture. There may well be a group of models that provide purely local computation, along with a means of determining when each model is applicable. Higher-order differential models, however, or low-order differential models that require too much a priori scene knowledge do not appear practicable. For any realistic model it seems inevitable that local processing must play an important role. Consequently, what

can we compute locally from the shading data? This is the question we shall now address.

4. Analysis of Local Computation

The relationship between surface orientation and image irradiance for a uniformly diffuse reflecting surface that is doubly curved is given by the surface-orientation-image-irradiance equations of Section 2.3. Parameter counting reveals that local image measurements are insufficient to specify surface orientation for the general case, but shape constraints can overcome these degrees of freedom. Since we wish to calculate surface shape locally, we consider the case in which we can assume a constant curvature over the small surface patch from which we draw information for the local calculation. Of course the curvature varies with direction; we only assume that we can ignore any change in curvature over the surface patch. Of course, this assumption is not valid in general; we are restricting our attention to this case to simplify the analysis. If we cannot determine what shape information is available in this restricted case, we are not likely to understand the general case. For this case, when we ignore curvature change, $l_{xx} = l_{yy} = l_{xy} = m_{xx} = m_{yy} = m_{xy} = 0$, and from the surface-orientation-image-irradiance equations we derive the expressions

$$\frac{I_{xx}}{I_{xy}} = \frac{l_x^2(1 - m^2) + m_x^2(1 - l^2) + 2l_x m_x l m}{l_x l_y(1 - m^2) + m_x m_y(1 - l^2) + (l_x m_y + l_y m_x) l m} ,$$

$$\frac{I_{yy}}{I_{xy}} = \frac{l_y^2(1 - m^2) + m_y^2(1 - l^2) + 2l_y m_y l m}{l_x l_y(1 - m^2) + m_x m_y(1 - l^2) + (l_x m_y + l_y m_x) l m} .$$

Notice that these relationships are only between surface shape and the second derivatives of the image irradiance. It is the assumption of constant curvature, not the more restrictive sphericity assumption (used by Pentland to recover surface orientation from the second derivatives of image irradiance), that is necessary to relate shape and just the second derivatives of the image irradiance. Image measurements are generally dependent on scene parameters other than those encoding shape. The first and second derivatives of image irradiance depend on the lighting position and the surface albedo, but the ratios of second derivatives are independent of all scene parameters except surface shape.

Can we use the above expressions to calculate surface orientation? We have previously [15] pointed to the insensitivity of surface-orientation-image-irradiance equations to surface orientation. The above expressions are also insensitive to surface orientation. We see this in the following considerations. Algebraic manipulation of the above expressions yields

$$\frac{I_{xx}}{I_{yy}} = \frac{l_x^2 + m_x^2 - (l_x m - l m_x)^2}{l_y^2 + m_y^2 - (l_y m - l m_y)^2} .$$

Suppose that over an image patch we know values of l and m that satisfy the above expression. Consider now this expression for $\frac{I_{xx}}{I_{yy}}$ when

$$l' = w_1 l \quad \text{and} \quad m' = w_2 m ,$$

at each point of the image patch. Using finite-difference formulas to calculate the derivatives of the surface normal, we obtain

$$\begin{aligned}\frac{I'_{xx}}{I'_{yy}} &= \frac{l_x'^2 + m_x'^2 - (l_x' m' - l' m_x')^2}{l_y'^2 + m_y'^2 - (l_y' m' - l' m_y')^2} \\ &= \frac{w_1^2 l_x^2 + w_2^2 m_x^2 - w_1^2 w_2^2 (l_x m - l m_x)^2}{w_1^2 l_y^2 + w_2^2 m_y^2 - w_1^2 w_2^2 (l_y m - l m_y)^2}\end{aligned}$$

Note that, as the magnitude of w_1 or w_2 is varied, the numerator and denominator of $\frac{I'_{xx}}{I'_{yy}}$ vary in like manner; both either increase or decrease; $\frac{I'_{xx}}{I'_{yy}}$ remains approximately equal to $\frac{I_{xx}}{I_{yy}}$. The ratios of the second derivatives of image irradiance are not sensitive to surface orientation. We cannot get further shape information from other image measurements, as the first and second derivatives of image irradiance are dependent on the surface albedo and the lighting conditions, and the image irradiance is dependent on surface albedo, lighting conditions, and the level of constant illumination from such sources as atmospheric scatter and the dark current of the sensor. Surface orientation can be computed locally only when very restrictive assumptions about surface shape are made. Without such restrictions there is not enough information in the shading to decouple surface orientation effects from those of albedo and illumination.

If shading is insufficient to allow surface orientation to be recovered, what then does the shading specify? Does it specify curvature? Can we compute it locally? Consider the above expressions for $\frac{I_{xx}}{I_{xy}}$, and $\frac{I_{yy}}{I_{xy}}$. Suppose that we know the correct values for l and m at an image point and we want to calculate l_x, l_y, m_x , and m_y . If l_x, l_y, m_x , and m_y is a solution, then so is $w l_x, w l_y, w m_x$, and $w m_y$, where w is any constant. Curvature cannot be computed locally (without further shape assumptions). The ratios of second derivatives of image irradiance contain shape information, yet are insensitive to surface orientation and do not allow computation of the curvature. What information about the surface do they encode?

To answer this question, we first rewrite the expressions for $\frac{I_{xx}}{I_{xy}}$ and $\frac{I_{yy}}{I_{xy}}$ in vector dot product form:

$$\begin{aligned}\frac{I_{xx}}{I_{xy}} &= \frac{[\frac{\partial}{\partial x}(l, m, \sqrt{1-l^2-m^2})] \cdot [\frac{\partial}{\partial x}(l, m, \sqrt{1-l^2-m^2})]}{[\frac{\partial}{\partial x}(l, m, \sqrt{1-l^2-m^2})] \cdot [\frac{\partial}{\partial y}(l, m, \sqrt{1-l^2-m^2})]} \\ \frac{I_{yy}}{I_{xy}} &= \frac{[\frac{\partial}{\partial y}(l, m, \sqrt{1-l^2-m^2})] \cdot [\frac{\partial}{\partial y}(l, m, \sqrt{1-l^2-m^2})]}{[\frac{\partial}{\partial x}(l, m, \sqrt{1-l^2-m^2})] \cdot [\frac{\partial}{\partial y}(l, m, \sqrt{1-l^2-m^2})]}\end{aligned}$$

Using the notation $\mathbf{N} = (l, m, \sqrt{1-l^2-m^2})$, for the unit surface normal, we obtain

$$\begin{aligned}\frac{I_{xx}}{I_{xy}} &= \frac{\mathbf{N}_x \cdot \mathbf{N}_x}{\mathbf{N}_x \cdot \mathbf{N}_y} \\ \frac{I_{yy}}{I_{xy}} &= \frac{\mathbf{N}_y \cdot \mathbf{N}_y}{\mathbf{N}_x \cdot \mathbf{N}_y}\end{aligned}$$

where $\mathbf{N}_x = \frac{\partial \mathbf{N}}{\partial x}$ and $\mathbf{N}_y = \frac{\partial \mathbf{N}}{\partial y}$.

For the case studied — when curvature changes are ignored — the ratios of the second derivatives of image irradiance measure the *relative squared curvature* of the surface. In other words, the ratios measure the relative change of the surface normal as we move along orthogonal image directions. However, relative curvature calculated locally at each image point constitutes insufficient information to allow surface shape reconstruction in the absence of further information about surface parameters. From shading information alone shape is an unattainable goal.

If we can find surface shapes, however, for which knowledge of relative curvature implies stronger information about the surface, e.g., surface orientation as in the case of a sphere, and if these surface shapes are reasonable approximations of the surfaces found in nature, then we may be able to recover stronger shape information locally. Locally there is not enough information to calculate surface shape without further knowledge, or without additional assumptions about surface shape. Pentland's work shows that an assumption of sphericity is strong enough to allow surface orientation to be calculated locally. Is this ability to calculate surface orientation specifically related to sphericity — or is it a feature that is generally true when we restrict the surface shape to cases in which the number of free parameters is no more than that for a spherical surface? In the foregoing discussion we have assumed that the surface is doubly curved. We shall now consider the images of singularly curved surfaces.

Just as we did for doubly curved surfaces, we assume that the derivatives of surface curvature can be ignored when we consider local computation of surface parameters. Differentiating the image irradiance equation, we obtain the same expression as before for the doubly curved surface, namely,

$$\frac{I_{xx}}{I_{yy}} = \frac{l_x^2 + m_x^2 - (l_x m - l m_x)^2}{l_y^2 + m_y^2 - (l_y m - l m_y)^2}.$$

For a singularly curved surface ($l_x m_y - l_y m_x = 0$) when surface curvature is locally constant, the second derivatives of image irradiance are not independent, $I_{xx} I_{yy} = I_{xy}^2$. Consequently, we can derive only one expression relating shape and the second derivatives of image irradiance, rather than the two expressions we derived for doubly curved surfaces. As before, it follows that

$$\frac{I_{xx}}{I_{yy}} = \frac{N_x \cdot N_x}{N_y \cdot N_y}.$$

At first, it might appear that there is more shape information in the first derivatives of image irradiance for

$$\begin{aligned} I_x m_y - I_y m_x &= 0, \\ I_x l_y - I_y l_x &= 0. \end{aligned}$$

But this is not the case, as the first and second derivatives of image irradiance are not independent. For singularly curved surfaces, when we ignore curvature change, $\frac{I_{xx}}{I_{yy}} = \left(\frac{I_x}{I_y}\right)^2$.

For the singularly curved and doubly curved surfaces studied, local shading specifies the relative curvature of the surface along orthogonal image coordinates, which is the most we can hope to recover by local computation. In general, we cannot ignore curvature change over a patch. In this case, the information available locally in the image combines data on

relative curvature and curvature change. In the restricted case in which the the surface is assumed to be spherical the surface orientation can be calculated. However, this appears to be a very special situation based on the sphericity assumption rather than on a restriction in the number of parameters needed to specify the surface. Since surfaces in general are not locally spherical, one is forced to conclude that shading alone cannot enable prediction of surface shape by purely local computation.

5. Conclusions

The recovery of a scene's surface shape from its image is fundamental to the vision process. Our purpose in processing an image is the recovery of scene properties, not those of the image per se. In remote sensing it is these scene properties that we wish to measure, but, to extract them, we have to understand how these scene properties are manifested in the image data. A conceptual model of the relationship between scene and image parameters is provided by intrinsic images. Each intrinsic image specifies, for each point in the image, the value of one of the scene parameters that contribute to the measured image intensity. Vision models try to recover these parameters as best they can, whereupon a type of relaxation process adjusts their values so that they constitute a consistent interpretation of the scene's structure. Which parameters are specified by separate intrinsic images and which are composite is unknown, but it is essential that they be estimable without the need to know the values of the other intrinsic images. Shape-from-shading proposes a source of information, namely shading, from which shape information is to be recovered — but what shape information does it actually encode?

Local shading specifies no more than the relative curvature of the scene's surface along orthogonal image directions. In general, even the recovery of relative curvature is complicated by change in the curvature of the surface. However, surface shape variables are related to image measurements in a fashion that is not dependent on knowing the other scene parameters. Shading provides direct shape information, but this is not enough for reconstruction of the surface shape. Further relationships between shape variables and image properties must be established before shape recovery is possible.

The various approaches reviewed have attempted to recover surface orientation from shading. To do so they have added extra information, such as known boundary conditions or constraints upon surface shape. The performance of these various models allows us to draw the following conclusions:

- Direct integration of differential models of scene properties requires much a priori information and has to contend with major computational problems.
- Local computation must play a major role in the recovery of scene parameters, but the models used have been overly restrictive in an effort to recover particular information.
- A relaxation mechanism, based on a strong low-order differential model, seems a viable means of propagating spatial information and constraints.

Shading provides a basis for an intrinsic image, specifying relative surface curvature and curvature change, but this intrinsic image alone is insufficient for surface shape recovery. Other models incorporating other image measurements are needed to complement shading.

Such models should utilize the advantages of local computation.

References

1. Barrow, H.G. and Tenenbaum, J.M., Recovering intrinsic scene characteristics from images, in: Hanson, A. and Riseman, E. (Eds.), *Computer Vision Systems*, Academic Press, New York, 1978 3-26.
2. Marr, D., Representing visual information, in: Hanson, A. and Riseman, E. (Eds.), *Computer Vision Systems*, Academic Press, New York, 1978 61-80.
3. Horn, B.K.P., Obtaining shape from shading information, in: Winston, P.H. (Ed.), *Psychology of Computer Vision*, McGraw-Hill, New York, 1975 115-155.
4. Horn, B.K.P., Understanding image intensities, *Artificial Intelligence* 8 (1977) 201-231.
5. Woodham, R.J., A cooperative algorithm for determining surface orientation from a single view, *Proceedings of the Fifth IJCAI*, Cambridge, Massachusetts, 1977 635-641.
6. Woodham, R.J., Relating properties of surface curvature to image intensity, *Proceedings of the Sixth IJCAI*, Tokyo, Japan, 1979 971-977.
7. Strat, T.M., A numerical method for shape from shading from a single image, S.M. Thesis, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology (1979).
8. Ikeuchi, K. and Horn, B.K.P., Numerical shape from shading and occluding boundaries, *Artificial Intelligence* 17 (1981) 141-184.
9. Pentland, A.P., The visual inference of shape: computation from local features, Ph.D. Thesis, Department of Psychology, Massachusetts Institute of Technology (1982).
10. Smith, G.B., From image irradiance to surface orientation, (submitted for publication). Available as Technical Note 273, Artificial Intelligence Center, SRI International (1982).
11. Horn, B.K.P. and Sjoberg R.W., Calculating the reflectance map, A.I. Memo 498, Artificial Intelligence Laboratory, Massachusetts Institute of Technology (1978).
12. Bruss, A.R., The image irradiance equation: its solution and application. AI-TR-623, Artificial Intelligence Laboratory, Massachusetts Institute of Technology (1981).
13. Brooks, M.J., Surface normals from closed paths, *Proceedings of the Sixth IJCAI*, Tokyo, Japan, 1979 98-101.
14. Brown, C.M., Ballard, D.H. and Kimhall, O.A., Constraint interaction in shape-from-shading algorithms, *Proceedings : Image Understanding Workshop*, Palo Alto, 1982 79-89.
15. Smith, G.B., The relationship between image irradiance and surface orientation, *Proceedings Computer Vision and Pattern Recognition Conference*, Washington, D.C. 1983.
16. Courant, R. and Hilbert, D., *Methods of Mathematical Physics*, Volume 2, Interscience Publishers, New York, 1962.

Appendix A

If a surface is twice differentiable, then

$$l_y(1 - m^2) + m_y lm = m_x(1 - l^2) + l_x lm \quad .$$

We call this the surface continuity equation, even though surface continuity is less demanding than the requirement that the surface be twice differentiable.

Proof: For a continuous twice-differential surface,

$$z_{xy} = z_{yx} \quad .$$

But $p = z_x$ and $q = z_y$, so

$$p_y = q_x \quad .$$

However,

$$p = \frac{-l}{\sqrt{1 - l^2 - m^2}} \quad ,$$

$$q = \frac{-m}{\sqrt{1 - l^2 - m^2}} \quad .$$

Hence,

$$p_y = -\frac{l_y(1 - m^2) + m_y lm}{(1 - l^2 - m^2)^{\frac{3}{2}}} \quad ,$$

$$q_x = -\frac{m_x(1 - l^2) + l_x lm}{(1 - l^2 - m^2)^{\frac{3}{2}}} \quad .$$

Then, substituting in $p_y = q_x$ yields

$$l_y(1 - m^2) + m_y lm = m_x(1 - l^2) + l_x lm \quad .$$

Appendix B

Developable surfaces are characterized by the differential equation

$$l_x m_y - l_y m_x = 0 \quad .$$

Proof: With the exception of a cylinder whose axis is parallel to the z axis, the differential equation defining all developable surfaces is [16]

$$z_{xx} z_{yy} - z_{xy}^2 = 0 \quad .$$

As the surface is twice differentiable, then $z_{xy} = z_{yx}$ so

$$z_{xx} z_{yy} - z_{xy} z_{yx} = 0 \quad .$$

As $p = z_x$ and $q = z_y$ then

$$p_x q_y - p_y q_x = 0 \quad .$$

But

$$p = \frac{-l}{\sqrt{1-l^2-m^2}} \quad ,$$

$$q = \frac{-m}{\sqrt{1-l^2-m^2}} \quad .$$

Hence,

$$p_x = -\frac{l_x(1-m^2) + m_x l m}{(1-l^2-m^2)^{\frac{3}{2}}} \quad ,$$

$$p_y = -\frac{l_y(1-m^2) + m_y l m}{(1-l^2-m^2)^{\frac{3}{2}}} \quad ,$$

$$q_x = -\frac{m_x(1-l^2) + l_x l m}{(1-l^2-m^2)^{\frac{3}{2}}} \quad ,$$

$$q_y = -\frac{m_y(1-l^2) + l_y l m}{(1-l^2-m^2)^{\frac{3}{2}}} \quad .$$

Substituting in $p_x q_y - p_y q_x = 0$ gives

$$l_x m_y - l_y m_x = 0 \quad .$$

Appendix C

The relationships between surface orientation and image irradiance for a developable surface are

$$I_x^2(m_y(1-l^2) + l_y l m) = I_y^2(l_x(1-m^2) + m_x l m) \quad ,$$

$$I_x m_y - I_y m_x = 0 \quad ,$$

$$I_x l_y - I_y l_x = 0 \quad .$$

Proof: Differentiating the image irradiance equation, $I(x, y) = R(l, m)$, we obtain

$$I_x = R_l l_x + R_m m_x \quad ,$$

$$I_y = R_l l_y + R_m m_y \quad .$$

Now

$$I_x^2(m_y(1-l^2) + l_y l m) = R_l^2 l_x (l_x m_y(1-l^2) + l_x l_y l m) \\ + R_m^2 m_x (m_x m_y(1-l^2) + l_y m_x l m) \\ + 2R_l R_m l_x (m_x m_y(1-l^2) + l_y m_x l m) \quad ,$$

$$I_y^2(l_x(1-m^2) + m_x l m) = R_l^2 l_y (l_x l_y(1-m^2) + l_y m_x l m) \\ + R_m^2 m_y (l_x m_y(1-m^2) + m_x m_y l m) \\ + 2R_l R_m l_y (l_x m_y(1-m^2) + m_x m_y l m) \quad .$$

But, for a developable surface $l_x m_y = l_y m_x$, (see Appendix B); hence

$$\begin{aligned} I_x^2(m_y(1-l^2) + l_y l m) &= R_l^2 l_x (l_y m_x(1-l^2) + l_x l_y l m) \\ &\quad + R_m^2 m_x (m_x m_y(1-l^2) + l_x m_y l m) \\ &\quad + 2R_l R_m l_x (m_x m_y(1-l^2) + l_x m_y l m) \quad , \\ I_y^2(l_x(1-m^2) + m_x l m) &= R_l^2 l_y (l_x l_y(1-m^2) + l_x m_y l m) \\ &\quad + R_m^2 m_y (l_y m_x(1-m^2) + m_x m_y l m) \\ &\quad + 2R_l R_m l_y (l_y m_x(1-m^2) + m_x m_y l m) \quad . \end{aligned}$$

Therefore,

$$\begin{aligned} I_x^2(m_y(1-l^2) + l_y l m) &= (R_l^2 l_x l_y + R_m^2 m_x m_y + 2R_l R_m l_x m_y)(m_x(1-l^2) + l_x l m) \quad , \\ I_y^2(l_x(1-m^2) + m_x l m) &= (R_l^2 l_x l_y + R_m^2 m_x m_y + 2R_l R_m l_y m_x)(l_y(1-m^2) + m_y l m) \quad . \end{aligned}$$

However, the surface continuity equation, (see Appendix A), is

$$l_y(1-m^2) + m_y l m = m_x(1-l^2) + l_x l m \quad .$$

We have the required result, i.e., that the relationship between surface orientation and image irradiance for a developable surface is

$$I_x^2(m_y(1-l^2) + l_y l m) = I_y^2(l_x(1-m^2) + m_x l m) \quad .$$

In terms of p and q , the equivalent form is

$$I_x^2 q_y - I_y^2 p_x = 0 \quad .$$

In terms of depth z , the equivalent form is

$$I_x^2 z_{yy} - I_y^2 z_{xx} = 0 \quad .$$

Note that, in addition,

$$\begin{aligned} I_x m_y - I_y m_x &= R_l(l_x m_y - l_y m_x) \quad , \\ I_x l_y - I_y l_x &= R_m(l_y m_x - l_x m_y) \quad . \end{aligned}$$

Hence, for a developable surface $l_x m_y - l_y m_x = 0$, we obtain the required results

$$\begin{aligned} I_x m_y - I_y m_x &= 0 \quad , \\ I_x l_y - I_y l_x &= 0 \quad . \end{aligned}$$